

EXHIBIT 3

UYLESS
BLACK

SON

COM

ONS

Uyless Black

TCP/IP and Related Protocols

- *ICMP*
- *FTP*
- *TELNET*
- *SMTP*
- *ARP*
- *RARP*
- *OSPF*
- *RIP*
- *EGP*
- *BOOTP*
- *UDP*
- *IGP*
- *SNMP*
- *CMOT*
- *Domain Name System*

Library of Congress Cataloging-in-Publication Data

Black, Uyless D.

TCP/IP and related protocols / Uyless Black.

p. cm. — (Uyless Black series on computer communications)

Includes index.

ISBN 0-07-005553-X

1. Computer networks. 2. Computer network protocols. I. Title.

II. Series.

TK5105.5.B5663 1992

004.6'2—dc20

91-37187

CIP

*McGraw-Hill Series on Computer Communications, Uyless Black,
Series Advisor*

Copyright © 1992 by McGraw-Hill, Inc. All rights reserved. Printed in the United States of America. Except as permitted under the United States Copyright Act of 1976, no part of this publication may be reproduced or distributed in any form or by any means, or stored in a data base or retrieval system, without the prior written permission of the publisher.

3 4 5 6 7 8 9 0 DOC/DOC 9 8 7 6 5 4 3 2

ISBN 0-07-005553-X

The sponsoring editor for this book was Neil Levine, the editing supervisor was Fred Dahl, and the production supervisor was Pamela A. Pelton. It was set in Century Schoolbook by McGraw-Hill's Professional Book Group composition unit.

Printed and bound by R. R. Donnelley & Sons Company.

Information contained in this work has been obtained by McGraw-Hill, Inc., from sources believed to be reliable. However, neither McGraw-Hill nor its authors guarantee the accuracy or completeness of any information published herein and neither McGraw-Hill nor its authors shall be responsible for any errors, omissions, or damages arising out of use of this information. This work is published with the understanding that McGraw-Hill and its authors are supplying information but are not attempting to render engineering or other professional services. If such services are required, the assistance of an appropriate professional should be sought.

6 Chapter One

Networks A, B, and C are often called subnetworks. The term does not mean that they provide fewer functions than a conventional network. Rather, it means that the three networks consist of a full logical network with the subnetworks contributing to the overall operations for internetworking. Stated another way, the subnetworks comprise an internetwork or an internet.

An internetworking gateway is designed to remain transparent to the end user application. Indeed, the end user application resides in the host machines connected to the networks; rarely are user applications placed in the gateway. This approach is attractive from several standpoints. First, the gateway need not burden itself with application layer protocols. Since they are not invoked at the gateway, the gateway can dedicate itself to fewer tasks, such as managing the traffic between networks. It is not concerned with application level functions such as database access, electronic mail, and file management.

Second, this approach allows the gateway to support any type of application because the gateway considers the application message as nothing more than a transparent protocol data unit (PDU).

In addition to application layer transparency, most designers attempt to keep the gateway transparent to the subnetworks and vice versa. That is, the gateway does not care what type of network is attached to it. The principal purpose of the gateway is to receive a PDU that contains adequate addressing information to enable the gateway to route the PDU to its final destination or to the next gateway. This feature is also attractive because it makes the gateway somewhat modular; it can be used on different types of networks.

However, it should be emphasized that this transparency is not achieved by magic. Software must be written to enable communications to take place between the subnetwork protocol and the gateway. These procedures are usually proprietary in nature, and standards do not describe this interface between the gateway and the subnetwork. The exception to this statement is the publication of IEEE, OSI, and Internet service definitions that describe procedures (in an abstract way) between the host and gateway protocols (layers). These service definitions are examined later in this book.

Connectionless and connection-oriented protocols

The concept of connectionless and connection-oriented operations is fundamental to any communications protocol, and the Internet standards use both. It is essential that the reader have a clear understanding of their features. Their principal characteristics are as follows:

- *Connection-oriented operations:* A user and a network set up a logical connection before the transfer of data occurs. Usually, some type of relationship is maintained between the data units being transferred through the user/network connection.
- *Connectionless-mode operations:* No logical connection between the user and the network is established prior to the data transmission. The data units are transmitted as independent units.

The connection-oriented service requires a three-way agreement between the two end users and the service provider (for instance, the network). It also allows the communicating parties to negotiate certain options and quality of service (QOS) functions. During the connection establishment, all three parties store information about each other, such as addresses and QOS features. Once data transfer begins, the protocol data units (PDUs) need not carry much overhead protocol control information (PCI). All that is needed is an abbreviated identifier to allow the parties to access the tables and look up the full addresses and QOS features. Since the session can be negotiated, the communicating parties need not have prior knowledge of all the characteristics of each other. If a requested service cannot be provided, any of the parties can negotiate the service to a lower level or reject the connection request.

The connection-oriented service also provides (with a few exceptions) for the acknowledgment of all data units. Additionally, if problems occur during the transmission, a connection-oriented protocol provides mechanisms for the retransmission of the errored units. In addition to these services, most connection-oriented protocols ensure that the data arrives in the proper order at the final destination. Figure 1.3 summarizes the characteristics of connection-oriented networks.

The connectionless type of service manages user PDUs as independent and separate entities. No relationship is maintained between successive data transfers, and few records are kept of the ongoing

- Connection mapped through network
- Abbreviated addressing
- Usually fixed routing between networks
- Accountability provided

Figure 1.3 Connection-Oriented Networks.

802.3, 802.4, and 802.5. The LLC includes 802.2. This sublayer was implemented to make the LLC sublayer independent of a specific LAN access method. The LLC sublayer is also used to provide an interface into or out of the specific MAC protocol.

The MAC/LLC split provides several attractive features. First, it controls access to the shared channel among the autonomous user devices. Second, it provides for a decentralized (peer-to-peer) scheme that reduces the LAN's susceptibility to errors. Third, it provides a more compatible interface with WANs, since LLC is a subset of the HDLC superset. Fourth, LLC is independent of a specific access method while MAC is protocol specific. This approach gives an 802 network a flexible interface with upper layer protocols (ULPs), such as IP or the OSI's connectionless network protocol (CLNP) (discussed in Chap. 5).

Classes of service

The 802 LAN standards include four types of service for LLC users:

- Type 1:* Unacknowledged connectionless service
- Type 2:* Connection-oriented service
- Type 3:* Acknowledged connectionless service
- Type 4:* All of the above services

All 802 networks must provide unacknowledged connectionless service (Type 1). Optionally, connection-oriented service can be provided (Type 2). Type 1 networks provide no ACKs, flow control, or error recovery. Type 2 networks provide connection management, ACKs, flow control, and error recovery. Type 3 networks provide no connection setup or disconnect, but they do provide for the acknowledgement of data units.

Most Type 1 networks use a higher-level protocol (i.e., TCP in the transport layer) to provide connection management functions. IP can rest over LLC as well. Therefore, a LAN layered model could be as follows: Physical, MAC, LLC, IP, TCP, and an application layer.

Chapter 5 discusses the relationship of TCP/IP and LLC in more detail.

Repeaters, Bridges, Routers, Brouters, and Gateways

Networks were originally conceived to be fairly small systems consisting of relatively few machines. As the need for data communication

services has grown, it has become necessary to connect networks together for the sharing of resources and distribution of functions and administrative control. In addition, some LANs, by virtue of their restricted distance, often need to be connected together through other devices. These devices are called a number of names in the industry; in this section we will explain and define each of these machines.

Figure 2.5 shows the relationships of these devices vis-à-vis a layered model. A *repeater* is used to connect the media on a LAN, typically called media segments. The repeater has no upper layer functions; its principal job is to terminate the signal on one LAN segment and regenerate it on another LAN segment.

The term *bridge* is usually associated with an internetworking unit (IWU). It operates at the data link layer (always at the MAC sublayer and sometimes at the LLC sublayer). Typically, it uses MAC physical addresses to perform its relaying functions. As a general rule, it is a fairly low-function device and connects networks that are homogeneous (for example, IEEE-based networks).

A *router* operates at the network layer because it uses network layer addresses (for example, IP, X.121, E.164 addresses). It usually contains more capabilities than a bridge and may offer flow control mechanisms as well as source routing or nonsource routing features (discussed in the next section).

The term *gateway* is used to describe an entity (a machine or software module) which not only performs routing capabilities but may act as a protocol conversion or mapping facility (also called a convergence function). For example, such a gateway could relay traffic and also provide conversion between two different types of mail transfer applications.

Yet another term that has entered the market is *brouter*. (As if there were not enough terms.) The term *brouter* is used to describe a machine that combines the features of a router and a bridge. At first

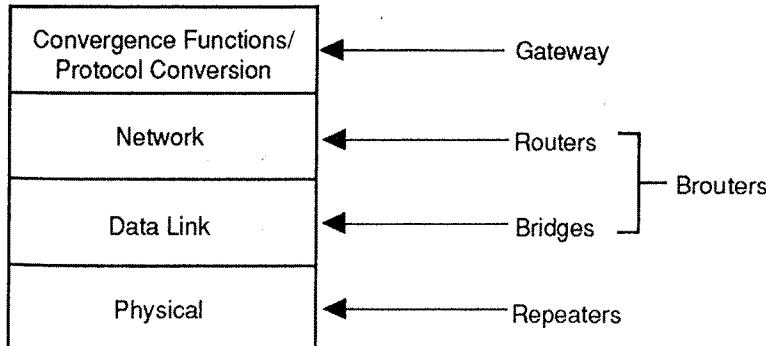


Figure 2.5 Placement of Internetworking Operations.

glance this seems redundant, but we shall see that the brouter is a powerful and flexible addition to internetworking products.

To avoid any confusion about these terms, some people use the term *internetworking unit* (IWU). An IWU is a generic term to describe a router, a gateway, a bridge, or anything else that performs relaying functions between networks.

Source routing and spanning tree bridges

The method in which internetworking protocol data units (PDUs) (datagrams or packets) are routed between networks is sometimes a source of confusion. The two major methods to perform routing are *source routing* and *nonsource routing*. Source routing derives its name from the fact that the transmitting device (the source) dictates the route of the PDU through an internet. The source (host) machine places the addresses of the “hops” (the intermediate networks or IWUs) in the PDU. Such an approach means that the internetworking units need not perform address maintenance, but they simply use an address in the routing field to determine where to route the frame.

In contrast, nonsource routing (using spanning tree techniques) makes decisions about the route and does not rely on the PDU to contain information about the route. Spanning tree routing is usually associated with nonsource routing and bridges and is quite prevalent in LANs.

The TCP/IP protocol suite utilizes source or nonsource routing but does not use spanning tree logic. These operations are found in the IP module and are introduced here and discussed in more detail in Chaps. 5 and 8.

An example of source routing on a LAN is illustrated in Fig. 2.6. The routing information field contains the LAN and bridge identifiers for each intermediate hop through the LAN network. Routing is accomplished by each bridge examining successive LAN numbers and bridge numbers in the routing information field and making a routing decision accordingly. As an example in Fig. 2.6, bridge 5 may receive a frame from LAN 3. Based on the routing information in the routing field, it might be required to route the frame out of its port to LAN 6 or out of another port to LAN 2. Again, under these conditions, the bridge has no control over how to route the frame.

Figure 2.7 depicts the operations of a spanning tree bridge. The bridge processor in the spanning tree bridge forwards frames based on an examination of the destination address. It compares this address to its bridge and routing information database. If the destination address is found in the forwarding table in this database, it determines the direction of the frame. If the frame is not intended for the port from which it came, it is forwarded on the correct port to the address indi-